

REMARKS

The present Amendment amends claims 1 and 13-15, and cancels claims 2-12 and 16. Therefore, the present application has pending claims 1 and 13-15.

Specification

The Examiner objected to the title as not being descriptive of the invention to which the claims are directed. Applicants have amended the title, as suggested by the Examiner. Therefore, this objection is overcome and should be withdrawn.

35 U.S.C. §102 Rejections

Claims 1-6, 8, 11-13, and 15-16 stand rejected under 35 U.S.C. §102(b) as being anticipated by U.S. Patent No. 5,987,477 to Schmuck, et al. ("Schmuck"). As indicated above, claims 2-6, 8, 11, and 12 were canceled. Therefore, this rejection with respect to claims 2-6, 8, 11, and 12 is rendered moot. This rejection with respect to the remaining claims 1, 13 and 15 is traversed for the following reasons. Applicants submit that the features of the present invention, as now more clearly recited in claim 1, 13 and 15, are not taught or suggested by Schmuck, whether taken individually or in combination with any of the other references of record. Therefore, Applicants respectfully request the Examiner to reconsider and withdraw this rejection.

Amendments were made to the claims to more clearly describe features of the present invention. Specifically, amendments were made to more clearly describe that the present invention is directed to a cache storage system as recited, for example, in independent claim 1.

The present invention, as recited in claim 1, provides a cache storage system connected to a plurality of clients and at least one storage device through a network. The cache storage system includes a controller and a device that stores data. The controller includes a lock management table containing lock status and an address correspondence table for indicating correspondence table that indicates correspondence between areas in the cache storage system and areas in the storage device. In a cache storage system, requests from clients for writing data may compete with each other. As such, the controller controls the device so as to temporarily store block data, which is exchanged between a client and the at least one storage device through the network, and which designates a logical address on a storage medium and a data length. The controller, upon receiving from the client a first lock request for an area of the device, checks the lock status of the area indicated in the address correspondence table and sends first status information of either conflict or good to the client. In addition, after sending the first status information of good to the client, the controller issues a first command of lock OK to the client to receive a first command of lock Acknowledge from the device, such that deadlock can be avoided when a first lock request is received from another client based on contents of the lock management table. After receiving the first command of lock Acknowledge and a first I/O request from the device, thereafter, the controller updates the contents of the lock management table to indicate first lock ON for the area of the device and processes data under the first I/O request and then issues a first I/O response to the client until the controller receives a first release command

from the client. In response to the first release command, the controller unlocks the area by updating the lock management table to indicate first lock OFF for the area of the device, and takes over processing to a channel control processor to control an I/O path connected to the storage device. The controller issues a second lock request to the at least one storage device to lock an area of the at least one storage device. After receiving either second status information of good or second status information of reservation conflict from the at least one storage device, and upon reception of a second command of lock OK, the controller sends a second command of lock acknowledge to the at least one storage device, and then updates the contents of the lock management table to indicate second lock ON for the area of the at least one storage device, so that deadlock can be avoided when a second lock request is received from another cache storage device or client. The controller also issues a second I/O request concerning updating of the data to the at least one storage device until the controller issues a second unlock request to the at least one storage device to unlock the area of the at least one storage device. Upon reception of a second command of unlock OK from the at least one storage device, the controller sends a second command of unlock Acknowledge to the at least one storage device to unlock the area of the at least one storage device. The prior art does not disclose all these features.

The above described features of the present invention, as now more clearly recited in the claims, are not taught or suggested by any of the references of record,

particularly Schmuck, whether taken individually or in combination with any of the other references of record.

Schmuck discloses a parallel file system and method for parallel write sharing. However, there is no teaching or suggestion in Schmuck of the cache storage system of the present invention, as recited in the claims.

Schmuck's computer system includes a shared disk file system running on multiple computers, each having its own instance of an operating system. Each computer is coupled for parallel data sharing access to files residing on network attached shared disks. The shared disk implementation, which is shown in Fig. 1, includes a token manager 11 that provides locking facilities for the computers at nodes 1, 2 and 3. Any available path can be used to avoid a server as a bottleneck or as a single point of failure. Since the required central functions in the lock manager have no attachment to a specific computer, they can be migrated from computer to computer to satisfy performance and availability needs.

One feature of the present invention, as recited in claim 1, includes where the controller controls the device so as to temporarily store block data, which is exchanged between a client and the at least one storage device through the network, and which designates a logical address on a storage medium and a data length. Schmuck does not disclose this feature. To support the assertion that Schmuck teaches this feature, the Examiner cites column 3, lines 26-49. However, neither the cited text nor any other portions of Schmuck disclose the claimed features. The cited text merely discloses a shared disk file system that supports disk

read and write calls with associated management calls. This is not the same as temporarily storing block data, which is exchanged between the client and the at least on storage device through a network, and which designates a logical address on a storage medium and a data length, as claimed.

Another feature of the present invention, as recited in claim 1, includes where the controller, upon receiving from the client a first lock request for an area of the device, checks the lock status of the area indicated in the address correspondence table and sends first status information of either conflict or good to the client. Schmuck does not disclose this feature. To support the assertion that Schmuck discloses an address correspondence table, the Examiner cites column 27, line 57 to column 28, line 10 (see rejection of claim 8 at pages 4-5 of Office Action). However, as described in the cited text, Schmuck merely describes the use of an attribute index to lookup attribute values. An attribute file, which is more fully described at column 26, line 64 to column 27, line 7, stores attribute data. The attribute data consists of one of two of a sequence of entries, including an attribute entry or a free space entry. The attribute entry contains the value of a particular attribute, and the free space entry marks free marks free space within the attribute file, i.e., space that can be reused the next time it is necessary to add new attribute entries to the attribute file. This attribute index of attribute values is quite different from the address correspondence table of the present invention, where the address correspondence table indicates correspondence of areas in the cache storage system with areas in the storage device. Furthermore, because Schmuck fails to

disclose the address correspondence table, as claimed, it follows that Schmuck fails to teach or suggest checking lock status of the area indicated in the address correspondence table, as now more clearly recited in claim 1. Even further, Schmuck fails to teach or suggest where the controller, upon receiving from the client a first lock request for an area of the device, checks the lock status and sends first status information of either conflict or good to the client, as in the present invention.

Yet another feature of the present invention, as recited in claim 1, includes where after sending the first status information of good to the client, the controller issues a first command of lock OK to the client to receive a first command of lock Acknowledge from the device, such that deadlock can be avoided when a first lock request is received from another client based on contents of the lock management table. Schmuck does not disclose this feature. For example, at column 17, lines 48 to 55, Schmuck discloses where the file system locks regions to synchronize access to a map. There is no disclosure in the cited text, or any other portion of Schmuck of where a controller issues a first command of lock OK to the client to avoid deadlock, in the manner claimed.

Another feature of the present invention, as recited in claim 1, includes where after receiving the first command of lock Acknowledge and a first I/O request from the device, thereafter, the controller updates the contents of the lock management table to indicate first lock ON for the area of the device and processes data under the first I/O request and then issues a first I/O response to the client until the controller

receives a first release command from the client. Schmuck does not disclose this feature. For example, at column 17, lines 48 to 55, Schmuck discloses where the file system locks regions to synchronize access to a map. However, there is no disclosure in Schmuck of where the controller updates the contents of the lock management table to indicate first lock ON for the area of the device and processes data under the first I/O request and then issues a first I/O response to the client until the controller receives a first release command from the client, in the manner claimed.

Even further, another feature of the present invention, as recited in claim 1, includes where in response to the first release command, the controller unlocks the area by updating the lock management table to indicate first lock OFF for the area of the device, and takes over processing to a channel control processor to control an I/O path connected to the storage device. Schmuck does not disclose this feature. As described above, Schmuck discloses at column 17, lines 48 to 55, where the file system locks regions to synchronize access to a map. However, this is not the same as the claimed features.

Yet another feature of the present invention, as recited in claim 1, includes where controller issues a second lock request to the at least one storage device to lock an area of the at least one storage device. After receiving either second status information of good or second status information of reservation conflict from the at least one storage device, and upon reception of a second command of lock OK, the controller sends a second command of lock acknowledge to the at least one storage

device, and then updates the contents of the lock management table to indicate second lock ON for the area of the at least one storage device, so that deadlock can be avoided when a second lock request is received from another cache storage device or client. Schmuck does not disclose this feature. Schmuck discloses at column 17, lines 48 to 55, where the file system locks regions to synchronize access to a map. However, this is not the same as the claimed features.

Still yet another feature of the present invention, as recited in claim 1, includes where the controller also issues a second I/O request concerning updating of the data to the at least one storage device until the controller issues a second unlock request to the at least one storage device to unlock the area of the at least one storage device. Schmuck does not disclose this feature. Again, Schmuck discloses at column 17, lines 48 to 55, where the file system locks regions to synchronize access to a map. However, this is not the same as the claimed features.

Another feature of the present invention, as recited in claim 1, includes where upon reception of a second command of unlock OK from the at least on storage device, the controller sends a second command of unlock Acknowledge to the at least one storage device to unlock the area of the at least one storage device. Schmuck does not disclose this feature. Again, Schmuck discloses at column 17, lines 48 to 55, where the file system locks regions to synchronize access to a map. However, this is not the same as the claimed features.

Therefore, Schmuck fails to teach or suggest "wherein requests for accessing to write data from said clients may compete with each other, and wherein said

controller controls said device to temporarily store block data which are exchanged between a client and said at least one storage device through said network and which designate a logical address on a storage medium and a data length” as recited in claim 1.

Furthermore, Schmuck fails to teach or suggest “wherein said controller, upon receiving from said client a first lock request for an area of said device, checks the lock status of said area indicated in said address correspondence table and sends first status information of either conflict or good to said client” as recited in claim 1.

Even further, Schmuck fails to teach or suggest “wherein said controller, after sending said first status information of good to said client, issues a first command of lock OK to said client in order to receive a first command of lock Acknowledge from said device so that deadlock can be avoided when a first lock request is received from another client on the basis of contents of said lock management table” as recited in claim 1.

Still even further, Schmuck fails to teach or suggest “wherein said controller updates the contents of said lock management table to indicate first lock ON for said area of said device and processes, after receiving said first command of lock Acknowledge and thereafter a first I/O request from said device, data under said first I/O request and thereafter issues a first I/O response to said client until said controller receives a first release command from said client” as recited in claim 1.

Yet even further, Schmuck fails to teach or suggestion “wherein said controller, in response to said first release command, unlocks said area by updating

said lock management table thereby to indicate first lock OFF for said area of said device, and takeovers processing to a channel control processor to control an I/O path connected to said storage device” as recited in claim 1.

Furthermore, Schmuck fails to teach or suggest “wherein said controller issues a second lock request to said at least one storage device to lock an area of said at least one storage device, and after receiving either second status information of good or second status information of reservation conflict from said at least one storage device, and upon reception of a second command of lock OK, sends a second command of lock acknowledge to said at least one storage device, and thereafter updates the contents of said lock management table to indicate second lock ON for said area of said at least one storage device, so that deadlock can be avoided when a second lock request is received from another cache storage device or client” as recited in claim 1.

Further, Schmuck fails to teach or suggest “wherein said controller issues a second I/O request concerning updating of said data to said at least one storage device until said controller issues a second unlock request to said at least one storage device to unlock the area of said at least one storage device” as recited in claim 1.

Even further, Schmuck fails to teach or suggest “wherein said controller, upon reception of a second command of unlock OK from said at least one storage device, sends a second command of unlock Acknowledge to said at least one storage device to unlock said area of said at least one storage device” as recited in claim 1.

Therefore, Schmuck fails to teach or suggest the features of the present invention, as now more clearly recited in the claims. Accordingly, reconsideration and withdrawal of the 35 U.S.C. §102(b) rejection of claims 1, 13 and 15 are respectfully requested.

The remaining references of record have been studied. Applicants submit that they do not supply any of the deficiencies noted above with respect to the references used in the rejection of claims 1, 13 and 15.

35 U.S.C. §103 Rejections

Claims 9 and 14 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Schmuck in view of U.S. Patent Application Publication No. 2003/0187860 to Holland. As indicated above, claim 9 was canceled. Therefore, this rejection with regard to claim 9 is rendered moot. This rejection with regard to the remaining claim 14 is traversed for the following reasons. Applicants submit that claim 14 is dependent on claim 1. Therefore, claim 14 is allowable for at least the reasons set forth above regarding independent claim 1.

Claims 7 and 10 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Schmuck in view of U.S. Patent No. 5,590,300 to Lautzenheiser. As indicated above, claims 7 and 10 were canceled. Therefore, this rejection with regard to claims 7 and 10 is rendered moot.


In view of the foregoing amendments and remarks, Applicants submit that claims 1 and 13-15 are in condition for allowance. Therefore, early allowance of claims 1 and 13-15 is respectfully requested.

U.S. Application No. 10/764,600

To the extent necessary, Applicants petition for an extension of time under 37 CFR 1.136. Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, or credit any overpayment of fees, to the deposit account of Mattingly, Stanger, Malur & Brundidge, P.C., Deposit Account No. 50-1417 (referencing attorney docket no. 500.44034X00).

Respectfully submitted,

MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C.



Donna K. Mason
Registration No. 45,962

DKM/sdb
(703) 684-1120